Prepared for submission to JCAP

Skewness and Kurtosis as Indicators of Non-Gaussianity in Galactic Foreground Maps

Assaf Ben-David[a,b] **, Sebastian von Hausegger**[b] **and Andrew D. Jackson**[a]

[a]Niels Bohr International Academy, The Niels Bohr Institute, University of Copenhagen, Blegdamsvej 17, DK-2100 Copenhagen Ø, Denmark

[b]Discovery Center, The Niels Bohr Institute, University of Copenhagen, Blegdamsvej 17, DK-2100 Copenhagen Ø, Denmark

E-mail: bendavid@nbi.dk, s.vonhausegger@nbi.dk

Abstract. Observational cosmology is entering an era in which high precision will be required in both measurement and data analysis. Accuracy, however, can only be achieved with a thorough understanding of potential sources of contamination from foreground effects. Our primary focus will be on non-Gaussian effects in foregrounds. This issue will be crucial for coming experiments to determine B-mode polarization. We propose a novel method for investigating a data set in terms of skewness and kurtosis in locally defined regions that collectively cover the entire sky. The method is demonstrated on two sky maps: (i) the SMICA map of Cosmic Microwave Background fluctuations provided by the Planck Collaboration and (ii) a version of the Haslam map at 408 MHz that describes synchrotron radiation. We find that skewness and kurtosis can be evaluated in combination to reveal local physical information. In the present case, we demonstrate that the statistical properties of both maps in small local regions are predominantly Gaussian. This result was expected for the SMICA map. It is surprising that it also applies for the Haslam map given its evident large scale non-Gaussianity. The approach described here has a generality and flexibility that should make it useful in a variety of astrophysical and cosmological contexts.

arXiv:1509.03100v2 [astro-ph.CO] 18 Nov 2015

Contents

1 Introduction **1**

2 Skewness, kurtosis and non-Gaussianity **3**
2.1 Definitions and General Properties 3
2.2 Correlated Data 4

3 Analysis of Sky Maps **6**
3.1 The SMICA Map 6
3.2 The Haslam Map 7

4 Discussion **11**

1 Introduction

The release of the Planck Cosmic Microwave Background (CMB) temperature fluctuations data [1] represents a landmark in observational cosmology. The standard cosmological model (ΛCDM) of an isotropic and homogeneous Universe, emerging after an inflationary era and today containing mainly dark energy and cold dark matter, has proven to be highly successful in describing the CMB data. Its six parameters have been remarkably well determined and constrained [2]. Although there may be more that can be learned from CMB temperature data, the community is now turning its attention to measurements of CMB polarization signals with the specific aim of searching for signs of primordial B-mode polarization. This task presents some major challenges. This signal, if it exists, will be extremely faint relative to foreground effects and instrumental noise. While our understanding of the properties of various Galactic foregrounds contaminating the CMB signals has been sufficient for the temperature analyses, a much better level of measurement, understanding and modelling will be required in order to clean the CMB polarization sky maps to the level necessary for primordial B-mode detection. It is therefore essential to study the characteristics of foreground sky maps in detail, including their statistical properties. Synchrotron radiation is one of the most important of these foregrounds, and its temperature power spectrum has been the subject of several studies, e.g. [3, 4]. Issues related to the statistical properties of this foreground remain to be investigated. (Regarding B-mode maps, see [5])

In this work we present a general approach for studying sky maps by investigating their local one-point distribution in small patches. This approach has the advantage of being blind to the global structure of the map. An obvious feature of foreground maps is that they reflect structures in our Galaxy. It is unreasonable to expect that its myriad distinct, specific and complex features can be regarded as a whole suitable for a global analysis with statistical tools. On the other hand, simple statistical measures can provide useful tools for analyzing sufficiently small local patches of these maps.[1] This is simply an example of the fact that the properties of data sets are often scale-dependent. By selecting an appropriate patch size, it is possible to focus on a specific scale of interest and thus to add flexibility to the analysis.

[1] These patches, which will be defined more precisely below, are locally defined non-overlapping regions that collectively cover the entire sky.

Working with patches also enables us to identify regions of the sky with anomalous properties such as non-Gaussian regions of the CMB map and Gaussian regions of foreground maps. In addition, the characterization of local foreground properties can be an important aid in the design of experiments that wish to focus on small regions of the sky with favorable local properties. Given the large scale structures expected in foreground maps, it is natural to consider the dimensionless moments (i.e., moments that are independent of the local mean and variance) one patch at a time. Thus, in this work we investigate the use of the skewness and kurtosis as statistical probes. Certainly, many other statistical tests can and should be considered in the study of foreground sky maps, and our focus on the skewness and kurtosis will serve as a simple example of one such test. We hope that a broader variety of statistical tests employing, e.g., Minkowski functionals [6] and the Kullback–Leibler divergence [7], will improve our understanding of the various Galactic foregrounds and that they will help pave the way to high-precision analyses of CMB polarization patterns.

We will demonstrate our method of analysis using two sky maps with quite different properties. Since our primary focus is the investigation of foreground maps, we will consider the 408 MHz radio map of Haslam et al. [8] in a recent reprocessed version [9]. This sky map, dominated by Galactic synchrotron emission, is used to model the synchrotron foreground contamination of the CMB signal at higher frequencies [10]. It is considered to be non-Gaussian in nature and serves as a typical example of a foreground component relevant for CMB analysis. While its power on various scales has been studied and is understood (see, e.g. [3]), our analysis of the one-point distribution and its Gaussianity can be performed 'blindly', without invoking knowledge of the physics behind Galactic synchrotron emission. Before considering the Haslam map, however, we will first perform a similar investigation of the Planck 2015 SMICA sky map of CMB temperature fluctuations [11]. This map and its predecessors have been studied in considerable detail in recent years, and its properties are well-understood. It is known in particular that it can be regarded as consistent with a realization of a statistically homogeneous and isotropic random Gaussian process. Since the main motivation for our work is the study of sky maps, which naturally have the topology of a sphere, our examples and discussions involving spatial correlations etc. will employ the language and notations relevant for a sphere. This is done without loss of generality or relevance to other types of data sets.

The structure of the paper is as follows. In Section 2 we discuss general properties of the skewness and kurtosis of a data set consisting of uncorrelated draws and the complications that arise when the draws are correlated. We also describe our method of analysis. In Section 3 we apply our method to the Planck 2015 SMICA map (expected to be Gaussian) and the Haslam map (expected to be far from Gaussian). Sample regions of the Haslam map will be shown to illustrate some of the features that our method can identify. These results will be discussed and conclusions drawn in Section 4.

2 Skewness, kurtosis and non-Gaussianity

2.1 Definitions and General Properties

Consider a data set consisting of n random draws of the variable x on a given distribution. The moments of this data set relative to the mean, $\overline{x}$, are given simply as

$$m_k = \frac{1}{n}\sum_i (x_i - \overline{x})^k, \tag{2.1}$$

where m_2 is the usual variance. For present purposes we will be concerned with the moments in various patches of the sky. Thus, the sum in Eq. (2.1) should extend over all pixels in the patch under consideration, and the mean $\overline{x}$ is similarly that of the patch. For $k > 2$, it is useful to adopt $m_2^{1/2}$ as a unit of length in order to obtain moments, $\mu_k = m_k/m_2^{k/2}$, that depend on the shape of the underlying distribution but are independent of its scale. If the distribution in question is Gaussian, $\mu_k = 0$ for odd k and $\mu_k = (k-1)!!$ for even k in the limit $n \to \infty$. For studying the non-Gaussianity of a general distribution, it is convenient to measure these moments relative to their Gaussian values. Thus, the skewness and excess kurtosis are defined respectively as

$$\gamma_1 = \mu_3 \quad \text{and} \quad \gamma_2 = \mu_4 - 3. \tag{2.2}$$

The skewness and excess kurtosis of the elements of our data set have distributions whose low moments are known for the special case of Gaussian distributions. Specifically, the mean and variance values of γ_1 are

$$\begin{aligned} \mathrm{E}[\gamma_1] &= 0 \\ \mathrm{Var}[\gamma_1] &= \frac{6(n-2)}{(n+1)(n+3)}, \end{aligned} \tag{2.3}$$

while for γ_2 we have

$$\begin{aligned} \mathrm{E}[\gamma_2] &= -\frac{6}{n+1} \\ \mathrm{Var}[\gamma_2] &= \frac{24n(n-2)(n-3)}{(n+1)^2(n+3)(n+5)}. \end{aligned} \tag{2.4}$$

Further, each of these distributions is known to be normal in the limit of large n.

Comparison of the distribution of skewness *or* excess kurtosis resulting from random draws on an unknown distribution with such known results can provide a test of non-Gaussianity. We will argue here that it is of potentially greater value to consider the distribution of skewness *and* kurtosis. Before addressing this point, however, it is important to note that the values of γ_1 and γ_2 resulting from n random draws on any given distribution are not independent. To see this, we consider a single random draw of any size, n, on an arbitrary distribution. Without loss of generality, we shift and rescale these numbers to ensure mean 0 and variance 1. Calculate the average value of the non-negative function $(x-a)^2(x-b)^2$ in terms of the moments m_3 and m_4 for this draw, and minimize the result with respect to the real parameters a and b. The result of these operations is the Pearson inequality [12],

$$\gamma_2 \geq \gamma_1^2 - 2. \tag{2.5}$$

Note that the equality is satisfied if and only if the distribution permits only the values $x = a$ and $x = b$. Thus, although stronger inequalities can be constructed for special classes of distributions (e.g., unimodal or symmetric unimodal distributions [13, 14]), there is no stronger general inequality involving skewness and kurtosis.

As stated above, the distributions of skewness and kurtosis are normal in the large-n limit for the special case of a Gaussian. Noting that their associated variances differ by a factor of 4 in this limit, Jarque and Bera [15] were led to study the properties of the combination $(\gamma_1^2 + \gamma_2^2/4)$ and to show that its asymptotic behaviour in the large-n limit is precisely that of a χ^2-distribution with two degrees of freedom. In principle, comparison with this familiar distribution offers a more accurate measure of Gaussianity than either the skewness or kurtosis alone. In practice, convergence to the asymptotic limit is slow. This has prompted the introduction of ad hoc transformations[16] to render the skewness and kurtosis "more Gaussian" with the aim of improving the accuracy of this combined measure. The fact that these transformations treat skewness and kurtosis as independent means that they cannot respect the Pearson inequality and serves as a reminder that their utility is greatest in the region of maximum probability ($\gamma_1 = 0$ and $\gamma_2 \to 0$).

With the aid of Eqs. (2.3) and (2.4), the skewness and kurtosis can be useful in deciding whether a given data set is the result of uncorrelated random draws on a Gaussian or a non-Gaussian distribution. They are not a priori suitable for treating systems such as the SMICA map of CMB temperature fluctuations where the expectation of uncorrelated random draws in harmonic space necessarily implies the existence of correlations in pixel space. We will address this problem in the next section.

2.2 Correlated Data

When attempting to analyze the distribution of pixel values of sky maps in small patches, we must deal with the complication that these pixels are expected to be correlated. Maps depicting mainly Galactic foreground emissions are correlated on scales corresponding to the sizes of structures in the Galaxy, which can vary greatly when projected on the sphere. According to the standard cosmological model, however, maps of CMB temperature fluctuations, such as the Planck SMICA map, are expected to be correlated in pixel space. Although the harmonic coefficients can be assumed to have been drawn uncorrelated in harmonic space, the transformation to the pixel domain then ensures that the pixels are correlated. Assuming a homogeneous and isotropic distribution based on the angular power spectrum, C_ℓ, the angular correlation function depends only on the separation angle, θ, between pixels and is given by

$$C(\theta) = \sum_\ell \frac{2\ell + 1}{4\pi} C_\ell P_\ell(\cos\theta), \tag{2.6}$$

where P_ℓ are the Legendre polynomials. We can then define the expected angular scale of the correlation, θ_c, as

$$\theta_c^2 = \frac{C(0)}{|C''(0)|} = \frac{2\sum_\ell (2\ell+1) C_\ell}{\sum_\ell \ell(\ell+1)(2\ell+1) C_\ell}. \tag{2.7}$$

For example, in this work we will consider the SMICA map smoothed with a Gaussian kernel of $20'$ FWHM, which allows us to use only scales up to $\ell_{\rm max} = 1000$. If we choose $C_\ell = g_\ell^2 C_\ell^{\Lambda{\rm CDM}}$, where $C_\ell^{\Lambda{\rm CDM}}$ is the Planck best-fit ΛCDM power spectrum [17] and g_ℓ is the smoothing kernel, Eq. (2.7) results in a correlation angle $\theta_c = 29'$. When considering data in a patch

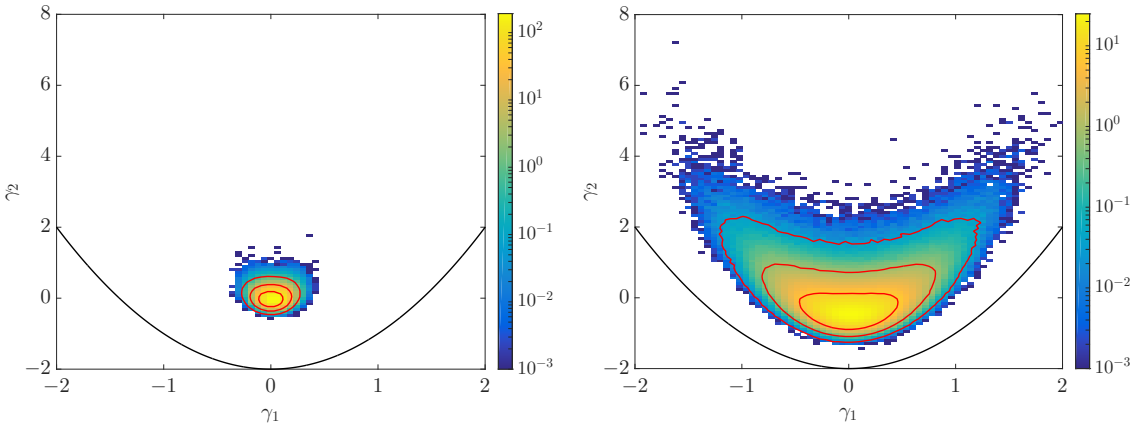

Figure 1. Mean histogram of skewness and excess kurtosis, calculated on (a) uncorrelated pixels and (b) correlated pixels with $\theta_\mathrm{p}/\theta_\mathrm{c} = 7.5$. The colors represent (mean) counts, on a logarithmic scale. Also plotted are the constraint parabola (*black*) imposed by Pearson's inequality (2.5) and the 1–3σ contour lines (*red*).

of angular size θ_p, we should expect that correlations will have a non-negligible effect on the joint distribution of the skewness and excess kurtosis of the data unless $\theta_\mathrm{p}/\theta_\mathrm{c} \gg 1$.

The effect of correlations is readily demonstrated. We use the hierarchical nature of HEALPix[2] [18] to partition a sky map of resolution $N_\mathrm{side} = 512$ into non-overlapping patches by assigning each pixel to a patch on a grid of $N_\mathrm{side} = 16$. This results in $N_\mathrm{p} = 3072$ patches with an angular size of $\theta_\mathrm{p} = 3.7°$, each containing 1024 pixels. We calculate the skewness and excess kurtosis of the pixels in each patch and bin them into a bivariate histogram. In order to smooth the statistical fluctuations of the histograms, we draw 1000 realizations of sky maps and calculate the mean of all histograms. We consider two such ensembles. The first is drawn in pixel space with pixels drawn uncorrelated from a normal distribution. The second ensemble is drawn in harmonic space, and we use the smoothed angular power spectrum C_ℓ mentioned above to draw the uncorrelated harmonic coefficients from a Gaussian distribution. This results in correlated pixels with $\theta_\mathrm{p}/\theta_\mathrm{c} = 7.5$. The two mean histograms are shown in Fig. 1. The difference between the two is striking. Pixel correlations cause the distributions of both γ_1 and γ_2 to be much wider, and their interdependence is made more apparent by the fact that the contour lines of the histogram are no longer elliptic.

We therefore require an alternative strategy for analyzing the joint distribution of γ_1 and γ_2 on the patches of a given sky map that will allow for correlations of the pixels. To this end, we have chosen a straightforward generalization of the strategy — common in one dimension — of comparing to an ensemble. Given an ensemble of realizations, we calculate the mean bivariate histogram, as those shown in Fig. 1, and use it as a proxy for the two dimensional probability distribution of the moments. This histogram, $h_{ij} = h(\gamma_{1,i}, \gamma_{2,j})$, where $\gamma_{1,i}$ and $\gamma_{2,j}$ are the bins of the moments, can now be used to assign a p-value to each point, (γ_1, γ_2), in the skewness–kurtosis plane in the following way. First, we calculate the value $h(\gamma_1, \gamma_2)$ by interpolation of the discrete histogram h_{ij}. We then find the bins comprising the appropriate 'tail' of the distribution, i.e. the set of all bins with smaller counts,

[2] http://healpix.sourceforge.net

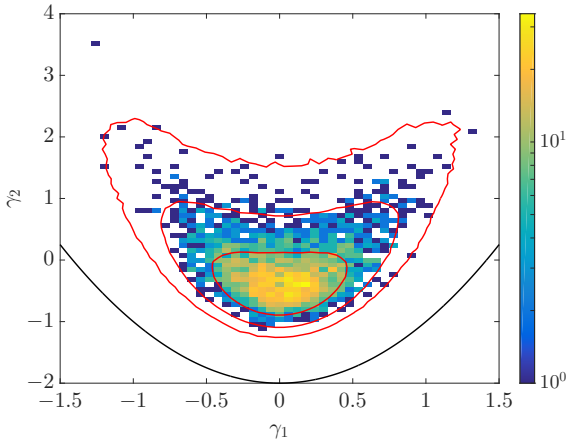

Figure 2. The skewness–kurtosis histogram for the SMICA map. Also plotted are the constraint parabola and the same contour lines as in Fig. 1(b).

$T = \{(i,j) \mid h_{ij} \leq h(\gamma_1, \gamma_2)\}$. Finally, the probability to exceed the given value is calculated as $p = N_{\rm p}^{-1} \sum_{(i,j)\in T} h_{ij}$, where the number of patches, $N_{\rm p}$, is obviously also the sum of h_{ij} over all bins. This procedure allows us to draw in Fig. 1 contours representing the 1–3σ bounds of the distribution, where we use the standard definition for the normal distribution to convert the language of p-values to that of σ-values, i.e. 1σ corresponds to $p = 31.7\%$, etc.

With a p-value assigned to each patch, it is possible to identify anomalous regions and to inspect them more closely. However, these p-values still depend on the nature of the ensemble of realizations chosen as a basis of comparison. We will utilize this procedure to analyze the SMICA and Haslam maps in the following section. We note that the assumption of statistical homogeneity and isotropy cannot be justified for foreground maps such as the Haslam map. While there is thus no guarantee that the strategy outlined here is quantitatively reliable when applied to foreground maps, we believe that it can provide useful qualitative guidance.

3 Analysis of Sky Maps

3.1 The SMICA Map

We start by analyzing the Planck 2015 SMICA map of CMB temperature fluctuations [11]. We degrade the map from its native HEALPix resolution of $N_{\rm side} = 2048$ to a resolution of $N_{\rm side} = 512$ after first smoothing it with a Gaussian kernel of $20'$ FWHM. Using an $N_{\rm side} = 16$ grid for the patches as described above, we obtain the skewness–kurtosis of the SMICA patches shown in Fig. 2. We note that the choice of smaller patches, e.g. by using a grid of $N_{\rm side} = 32$, would result in larger statistical fluctuations of the moments since each patch would contain fewer pixels. In addition, the smaller patch would be closer to the correlation angle of the map. While smaller patches could still be used for SMICA, they would not be appropriate for the Haslam map to be analyzed below.

In order to assess the expected distribution of the moments for the SMICA map, we draw a Gaussian ensemble for purposes of comparison. While the SMICA map also contains other effects (such as instrumental noise and foreground residuals) in addition to the dominant Gaussian signal, it was recently shown [7] that the statistical properties of a Gaussian ensemble are close to those of the Planck full focal plane (FFP) simulations, which include some of these additional effects on sufficiently large scales. We therefore use a Gaussian ensemble here and

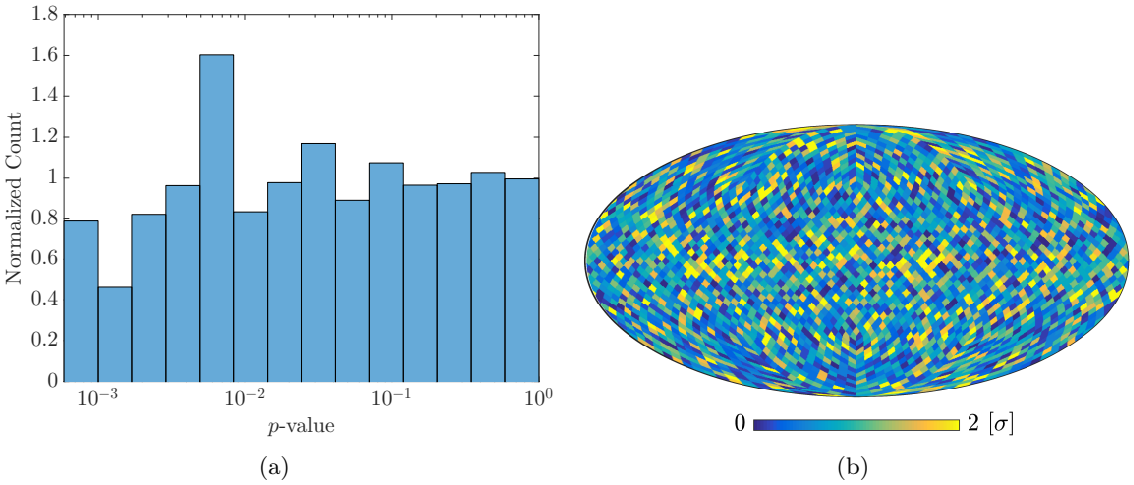

(a) (b)

Figure 3. (a) The distribution of p-values for the SMICA patches. The bin edges are logarithmically spaced and the histogram is normalized to unit total area. (b) The σ-level for each patch of the SMICA map, in Galactic coordinates.

will show that the statistical properties of the skewness and excess kurtosis of the SMICA map are consistent with it. Since we wish to ensure that the Gaussian ensemble has the same correlation structure in pixel space as the SMICA map, we use the ensemble of 1000 realizations based on the best-fit ΛCDM angular power spectrum described in Section 2.2. The mean histogram for this ensemble was shown in Fig. 1(b), and the 1–3σ contour lines calculated based on it are superimposed on the SMICA histogram in Fig. 2.

The contour lines strongly suggest that the distribution of patch p-values for the SMICA map follows the expected distribution, with $\approx 0.1\%$ of the patches crossing the 3σ boundary. We use the mean histogram to assign a p-value to each patch as described above and plot in Fig. 3(a) the distribution of the values.

It is apparent that the distribution is quite uniform. This provides confirmation that the SMICA map is indeed consistent with our ensemble of realizations and provides further support for the conclusion of [7] that the more complicated FFP simulations are not essential for analyzing the statistical properties of the SMICA map; simple Gaussian appear to be sufficient. Moreover, in the context of the present analysis, this conclusion can be extended to smaller scales with $\ell_{\rm max}$ as large as 1000.

Having labeled each patch with a p-value, or the corresponding σ-levels, we can now consider the spatial distribution of these values. We therefore plot in Fig. 3(b) a map showing the p-values, translated to σ-levels, of the patches of the SMICA map. Visual inspection of the SMICA patches suggests a tendency towards an excess of outliers in the Galactic plane, but we have not investigated this question quantitatively.

3.2 The Haslam Map

In addition to the SMICA map of CMB temperature fluctuations, we also apply our method to a reprocessed version [9] of the 408 MHz radio map of Haslam et al. [8]. We adopt this map as typical of the maps used as templates for cleaning foreground contributions to the CMB data. Understanding the statistical properties of such maps, or at least finding outlying regions in them, could improve the accuracy of CMB extraction. Since it is primarily a map of Galactic emissions, we do not expect the Haslam map as a whole to be Gaussian-distributed.

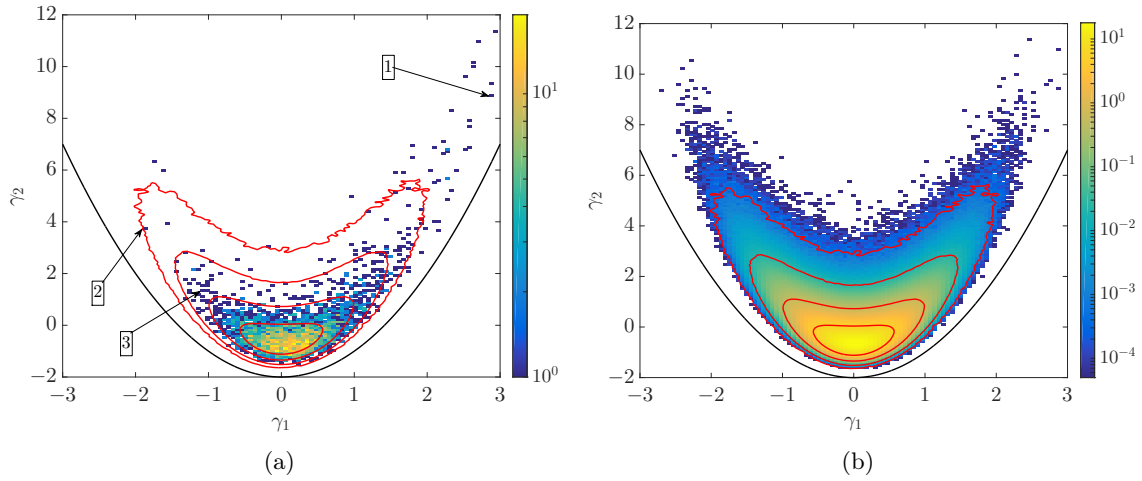

Figure 4. (a) The skewness–kurtosis histogram for the Haslam map. The bins of the three patches given as examples below are marked with arrows. (b) Mean histogram based on Gaussian realizations using the Haslam power spectrum, as explained in the text. The same set of 1–4σ contour levels, calculated using the mean histogram in (b), is shown in both panels (*red*).

Working in small patches, however, allows us to inspect the distributions in local areas of the map that may or may not be Gaussian. The $56'$ resolution of the map enables us to choose a patch grid (with $N_{\rm side} = 16$ but not smaller) identical to that chosen for the SMICA map. Since the Haslam map is provided in an $N_{\rm side} = 512$ pixelization, we are again left with 1024 pixels in each patch.

The skewness–kurtosis histogram for the Haslam map is shown in Fig. 4(a). The range of values of γ_1 and γ_2 is far larger than that found in the case of the SMICA map (Fig. 2). It is also immediately apparent that it is much more probable for a Haslam patch to have positive skewness than negative. We will comment more on this asymmetry below.

As before, we require an ensemble in order to assign a p-value to each patch. In contrast to the map of CMB fluctuations, however, the Haslam map is not believed to be a Gaussian realization drawn in harmonic space using some angular power spectrum. Thus, it is not clear how to generate a suitable ensemble. We note first that the ensemble should be Gaussian. This is a result of *our* wish to classify the patches by their level of Gaussianity and is unrelated to the intrinsic distribution of the data. The determination of the nature of this Gaussian or, equivalently, of the corresponding correlation matrix of the Haslam map would require detailed knowledge of the specific physical processes that govern synchrotron emissions in the Galaxy. In the interests of simplicity and generality, we prefer using the correlation structure of the Haslam map itself as a reference point. In order to obtain a statistical measure of the correlations, we are then led to treat the map as a realization of a homogeneous and isotropic distribution. We transform the map to the spherical harmonic domain and calculate the observed power spectrum,

$$\widehat{C}_\ell = \frac{1}{2\ell+1} \sum_{m=-\ell}^{\ell} |a_{\ell m}|^2, \tag{3.1}$$

where the $a_{\ell m}$ are the harmonic coefficients. We then use this power spectrum to draw Gaussian realizations in harmonic space. This results in an ensemble of maps having, on

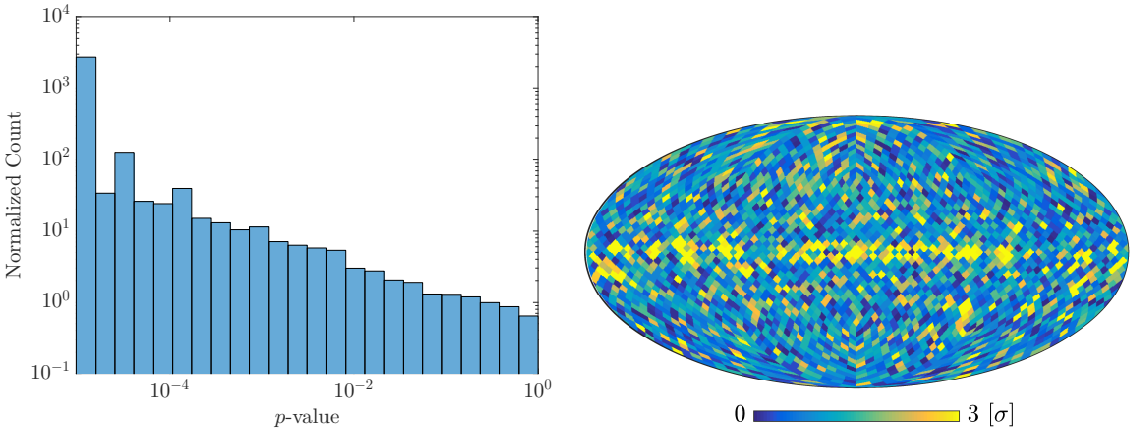

(a) (b)

Figure 5. (a) The distribution of p-values for the Haslam patches. The bin edges are logarithmically spaced and the histogram is normalized to unit total area. We note that the leftmost column here represents the sum of all patches with $p \lesssim 10^{-5}$. (b) The σ-level for each patch of the Haslam map, in Galactic coordinates.

average, the same structure of correlations as the Haslam map. We take $\ell_{\rm max} = 600$ since the Haslam angular power is negligible for smaller scales (larger values of ℓ). We draw 20 000 realizations, which allows us to reach a significance level of 4σ.

The mean histogram obtained using this ensemble is shown in Fig. 4(b). We use it to calculate the 1–4σ contour levels and superimpose them, as an illustration, on both histograms of Fig. 4. It is clear that many patches of the Haslam map are located far beyond the 4σ contour and are therefore potentially interesting. However, it is also intriguing that so many of the patches are located within the 1σ contour in light of the fact that the map as a whole is highly non-Gaussian. This is an illustration of the advantage of an analysis of small patches; it can highlight local features while disregarding global features of the map. We can now assign p-values to the patches. Their distribution is plotted in Fig. 5(a). We see that the distribution is not uniform. This indicates, as expected, that the Haslam map as a whole is not a consistent realization of the ensemble.

As with the SMICA map, we can examine the spatial distribution of the patches labeled by their significance levels. This map is shown in Fig. 5(b). It is immediately apparent that many of the highly non-Gaussian patches are located in the area of the Galactic plane. More surprisingly, it is evident that many of the patches in the Galactic plane are labeled with a very low value of σ. (We will suggest below that this result does not necessarily indicate that these patches are free of non-Gaussian foreground effects.) We also notice some highly non-Gaussian patches at high Galactic latitudes far from the Galactic plane.

The intriguing results of Fig. 5(b) provide motivation to examine some of the patches by eye in the hope of determining why they are anomalous. Our labeling of the patches by p-values (or σ-levels) provides a natural ordering scheme for their examination. We select three patches that are each characteristic of their locations in the skewness–kurtosis histogram, Fig. 4(a). Their positions in this histogram are indicated by the labels '1', '2' and '3'. In addition, in Fig. 6 we show the Haslam map and indicate on it the location of the patches. They lie in different areas of the map and have distinct characteristics. The panels of Fig. 7

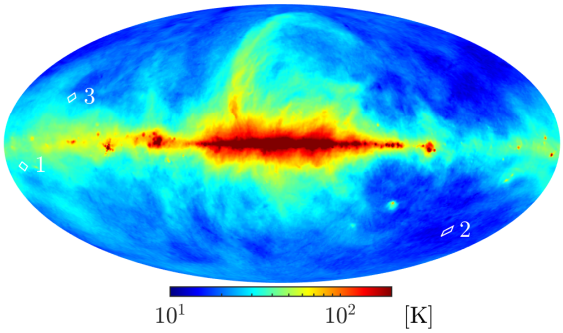

Figure 6. The reprocessed 408 MHz Haslam map, in Galactic coordinates. The locations of the three sample patches are marked in white.

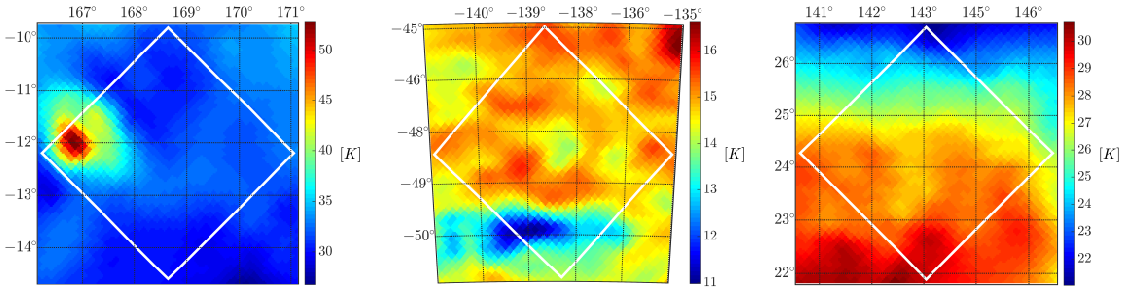

(a) (b) (c)

Figure 7. From left to right, patches 1, 2, and 3 and portions of their neighboring patches. The vertical and horizontal axes show Galactic latitude and longitude, respectively. The skewness and kurtosis of each patch can be seen in Fig. 4(a), and Fig. 6 shows their locations on the Haslam map.

show the patches in question (enclosed in white) along with portions of their neighboring patches. The first of these patches, Fig. 7(a), belongs to the relatively well-populated branch of the histogram with high kurtosis and large positive skewness. Such values arise naturally when the distribution of the patch contains a small area of relatively high intensity. These conditions are easily produced by a point source as is the case in the figure shown. The region of high intensity covers only a small fraction of the patch; the bulk has an intensity that is relatively low compared to the mean intensity of the patch. Patches with similar distributions are most frequently found along the Galactic plane. This fact is clearly visible in Fig. 5(b) and incidentally provides qualitative confirmation of the success of the method in [9], which was designed to remove extragalactic point sources from the original Haslam map.

We now consider patch 2—a patch selected from the less populated branch of the histogram in Fig. 4(a). The large magnitudes of the kurtosis and skewness again ensure that the p-value of this patch is small. In this case, however, the skewness is strongly negative. The patch should contain a small region of very low intensity while the complementary area should have an intensity somewhat above the average value. The obvious analogy to Fig. 7(a) thus makes it reasonable to expect to find a cold spot or point sink. Clearly, Fig. 7(b) meets this expectation. Although a cursory glance at the night sky is sufficient to remind us of the existence of point sources, it is more challenging to think of explanations for the phenomenon of cold spots. There are several possible explanations for the existence of regions where the intensity of synchrotron radiation is low. Small magnetic fields and/or low electron densities

can reduce the intensity of radiation emitted from a given region. Alternatively, intervening matter can absorb synchrotron radiation once it has been emitted. These effects can work individually or in concert. The relative paucity of patches with large negative skewness, evident in the histogram of Fig. 4(a), suggests that localized radiation sinks are significantly less likely than bright point sources.

The results reported here were obtained from a single patching pattern. It can happen, however, that the specific values of skewness and kurtosis of a given patch can be highly sensitive to its particular position on the map. Thus, a patch with high skewness and kurtosis need not necessarily contain a point source but can rather lie at an advantageous location, grazing a larger-scale structure. This can also occur for other values (γ_1, γ_2), e.g., also simulating cold point sources. Thus, patch 3 in Fig. 7(c) illustrates a case where the patch lies on top of a larger structure in which the temperature changes smoothly over the area of the patch. The value of the skewness—including its sign—will depend on the precise location of the patch relative to this fixed large scale structure.

The last case, not shown in a figure here, is that of random fluctuations. Large fractions of the sky are dominated by diffuse emission, which in the case of the Haslam map is closely connected to the morphology and fluctuations in the Galactic magnetic field and the local electron density. The pixels contained in patches lying in these regions will be distributed in a largely Gaussian manner. The histogram, Fig. 4(a), reveals that approximately 52% of the patches in the Haslam map are Gaussian with $\sigma \leq 1$, and inspection of Fig. 5(b) provides visual confirmation. The dominance of Gaussian patches was expected for the SMICA map and is quantitatively confirmed by Figs. 2–3. A similar dominance in the Haslam map was *not* expected. We will return to this surprising result in the following section.

4 Discussion

The results presented here are intended to illustrate the role that higher moments can play in understanding the statistical behaviour of foreground maps. In practice, it is necessary to adjust the size of the patches to ensure that two obvious criteria are met: The patches must be large on a scale set by the correlation angle, θ_c, in order to ensure that individual pixels have a reasonable degree of statistical independence. Second, the patches must be small in comparison with the size of diffuse foreground effects. Such considerations led us to choose patches with $N_{\rm side} = 16$. Clearly, patch size can be tuned to the scale of phenomena that we would like to identify. For example, the southern hemisphere cold spot seen in both WMAP and CMB temperature maps has an angular size of approximately 5°. As a consequence, it does not appear to be exceptional in the SMICA plot of Fig. 3(b) that was obtained with a patch size of 3.7°. It should be easy to spot with a larger patch size. We have also restricted our attention to a single realization of the "tiling" of the full sky. Since we have noted some phenomena that can be sensitive to the precise location of a patch, it would be prudent to consider several tilings displaced by an amount comparable to the patch size. Patch selection here was made using the HEALPix package that produces a subdivision of a spherical surface into patches of equal area, but our method can evidently be used for other patch forms and sizes. We also note that, although the decision to consider the statistics of skewness and kurtosis was natural, the same approach could be expanded to include a larger number of higher moments. The only technical challenge in such an extension would be the determination of the inequalities analogous to Eq. (2.5) that such higher moments must satisfy.

Turning to the results of our analysis, we note that the SMICA map of CMB temperature fluctuations is entirely consistent with the expectation of Gaussian behaviour. This agreement is confirmed with considerable accuracy by the fact that the distribution of p-values shown in Fig. 3(a) is independent of p.[3]

With a far larger fraction of patches with small p-values, the Haslam map would appear to confirm the a priori belief that this foreground map is strongly non-Gaussian. Closer inspection suggests that this conclusion may be premature. Figs. 7(a)–7(c) show that classification of a patch by its p-value alone is not sufficient to describe its full physical content. These figures reveal three classes of anomalous patches that can be distinguished by their structure in pixel space: point sources, cold spots and strong intensity gradients. Although it was not our goal to identify such structures, our method has the potential to do so mechanically. In practice, implementation should include shifting the location and changing the size of patches in order to avoid misidentification. Such localized phenomena differ distinctly from the diffuse effects of synchrotron radiation and merit special treatment when sky maps are to be cleaned. Their removal from Fig. 4(a) would eliminate many of the patches with $\sigma > 3$ and virtually all patches with $\sigma > 4$. The resulting histogram would have a significantly greater resemblance to a Gaussian result such as that shown in Fig. 4(b). There are two distinct explanations for this somewhat unexpected result. First, we note that Figs. 3(b) and 5(b) contain a large number of patches with $p \approx 1$ in the Galactic plane in spite of the fact that this region should be filled with sources of foreground effects. Mertsch and Sarkar [3] have argued that many independent foreground contributions can combine to create Gaussian distributions as a consequence of the central limit theorem.[4] For high Galactic latitudes, it is probably of greater importance to reconsider the assumptions of homogeneity and isotropy that led to the Gaussian realizations of the Haslam map shown in Fig. 4(b). While these assumptions are not expected to be valid for the map as a whole, they may well be valid locally on the relatively small scale of the patch size. In other words, large scale diffuse foreground effects can be described by correlations between spherical harmonics of relatively low order, $\ell \leq \ell_0$ without the involvement of components with $\ell > \ell_0$. These correlated terms will be approximately constant over patches whose characteristic size is smaller than $1/\ell_0$. They will not contribute to dimensionless moments such as the skewness and kurtosis that will be determined by the uncorrelated harmonics with large ℓ. In either language, no deeper explanation is required to understand the abundance of Gaussian patches.

It is of some interest to consider the consequences of the present results for a single portion of the sky. We have looked at the patches of the Haslam map that are associated with the region explored by the BICEP2 collaboration. We have sorted these patches according to their significance in six bins of width $\sigma/2$ ranging from 0 to 3σ. Since the edges of the BICEP2 zone are somewhat loosely defined, we have considered both a larger and a smaller zone. In the former case we considered all 58 patches that touch the BICEP2 zone and found the distribution $[14, 17, 13, 6, 5, 3]$ which can be compared with the Gaussian expectation of $[22.2, 17.4, 10.6, 5.1, 1.9, 0.5]$. The small zone consists of those 20 patches that lie completely within the BICEP2 zone. In this case the distribution is $[8, 6, 4, 2, 0, 0]$ which can be compared with the Gaussian result of $[7.7, 6.0, 3.7, 1.8, 0.7, 0]$. In both cases, there is good

[3] Bear in mind, for example, that the number of patches contributing to this histogram decreases rapidly as $p \to 0$.

[4] While this may not be the explanation of the global abundance of Gaussian patches seen in the Haslam map, it provides a useful reminder that the absence of evidence for the existence of non-Gaussian foreground effects does *not* constitute evidence of their absence.

agreement with the Gaussian expectations. Since the coupling between the foreground effects of synchrotron radiation and dust emission are expected to be large, there would appear to be grounds for the optimistic hope that the dominate foreground effects in the BICEP2 zone are actually Gaussian.

CMB science is now entering an era of high precision polarization analysis that will require a more precise foreground separation than was needed for the temperature analysis. We believe that the results presented here demonstrate that a patch-by-patch statistical study of the scale-free higher moments of the one-point distribution function can make a significant contribution to this program.

Acknowledgments

This work is based on observations obtained with Planck, an ESA science mission with instruments and contributions directly funded by ESA Member States, NASA, and Canada. We gratefully acknowledge the financial support of Danmarks Grundforskningsfond and of the Villum Foundation. It is a particular pleasure to thank Pavel Naselsky for the benefit of his invaluable advice and for his enthusiastic support.

References

[1] R. Adam *et al.* (Planck Collaboration), "Planck 2015 results. I. Overview of products and scientific results," (2015), submitted to A&A, arXiv:1502.01582 [astro-ph.CO] .

[2] P. A. R. Ade *et al.* (Planck Collaboration), "Planck 2015 results. XIII. Cosmological parameters," (2015), submitted to A&A, arXiv:1502.01589 [astro-ph.CO] .

[3] P. Mertsch and S. Sarkar, JCAP **1306**, 041 (2013), arXiv:1304.1078 [astro-ph.GA] .

[4] M. Regis, Astropart. Phys. **35**, 170 (2011), arXiv:1101.5524 [astro-ph.HE] .

[5] M. Kamionkowski and E. D. Kovetz, Phys. Rev. Lett. **113**, 191303 (2014).

[6] P. A. R. Ade *et al.* (Planck Collaboration), "Planck 2015 results. XVI. Isotropy and statistics of the CMB," (2015), submitted to A&A, arXiv:1506.07135 [astro-ph.CO] .

[7] A. Ben-David, H. Liu, and A. D. Jackson, JCAP **1506**, 051 (2015), arXiv:1506.07724 [astro-ph.CO] .

[8] C. G. T. Haslam, C. J. Salter, H. Stoffel, and W. E. Wilson, Astron. Astrophys. Suppl. Ser. **47**, 1 (1982).

[9] M. Remazeilles, C. Dickinson, A. J. Banday, M. A. Bigot-Sazy, and T. Ghosh, Mon. Not. R. Astron. Soc. **451**, 4311 (2015), arXiv:1411.3628 [astro-ph.IM] .

[10] R. Adam *et al.* (Planck Collaboration), "Planck 2015 results. X. Diffuse component separation: Foreground maps," (2015), submitted to A&A, arXiv:1502.01588 [astro-ph.CO] .

[11] R. Adam *et al.* (Planck Collaboration), "Planck 2015 results. IX. Diffuse component separation: CMB maps," (2015), submitted to A&A, arXiv:1502.05956 [astro-ph.CO] .

[12] K. Pearson, Phil. Trans. R. Soc. A **216**, 429 (1916).

[13] V. K. Rohatgi and G. J. Székely, Statistics & Probability Letters **8**, 297 (1989).

[14] C. A. Klaassen, P. J. Mokveld, and B. van Es, Statistics & Probability Letters **50**, 131 (2000).

[15] C. M. Jarque and A. K. Bera, Economics Letters **6**, 255 (1980).

[16] R. B. D'Agostino, Biometrika **57**, pp. 679 (1970).

[17] N. Aghanim *et al.* (Planck Collaboration), "Planck 2015 results. XI. CMB power spectra, likelihoods, and robustness of parameters," (2015), submitted to A&A, arXiv:1507.02704 [astro-ph.CO] .

[18] K. M. Gorski, E. Hivon, A. J. Banday, B. D. Wandelt, F. K. Hansen, M. Reinecke, and M. Bartelman, Astrophys. J. **622**, 759 (2005), arXiv:astro-ph/0409513 [astro-ph] .